

Towards an optimal decision strategy of visual search

Bo Chen, Pietro Perona*

Computations and Neural Systems, California Institute of Technology, Pasadena, CA, USA

* E-mail: perona@caltech.edu

Abstract

Searching for objects amongst clutter is a key ability of visual systems. Speed and accuracy are often crucial: how can the visual system trade off these competing quantities for optimal performance in different tasks? How does the trade-off depend on target appearance and scene complexity? We show that the optimal tradeoff strategy may be cast as the solution to a partially observable Markov decision process (POMDP) and computed by a dynamic programming procedure. However, this procedure is computationally intensive when the visual scene becomes too cluttered. Therefore, we also conjecture an optimal strategy that scales to large number of clutters. Our conjecture applies to homogeneous visual search and for a special case of heterogeneous search where the signal-to-noise ratio differs across location. Using the conjecture we show that two existing decision mechanisms for analyzing human data, namely diffusion-to-bound and maximum-of-output, are sub-optimal; the optimal strategy instead employs two scaled diffusions.

Introduction

One of the most useful functions of the visual system is searching for things: food, mates, threats. This is a difficult task: the relevant objects, whose appearance may not be entirely known in advance, are often embedded in irrelevant clutter. Furthermore, time is of the essence: the ability to detect quickly objects of interest is an evolutionary advantage. Speed comes at the cost of making mistakes. Shorter decision times imply collecting less signal, and expose the animal to detection errors. Thus, it is critical that each piece of sensory information is used efficiently to produce a decision in the shortest amount of time while maintaining the probability of errors within an acceptable limit. However, an ideal observer that describes the optimal trade-off between visual search speed (or response time (RT)) and error rate (ER) has not yet been proposed. By “optimal” we mean that the observer achieves the lowest expected RT among any alternative model with the same ER. Equivalently an ideal observer is the most accurate among all other models of the same speed.

Current visual search models fall into two categories. The first category is phenomenological, e.g. diffusion-to-bounds [1] and competitive accumulators [2]. Such models characterize well the RT versus ER trade-off in humans, but shed little light on optimality. Models of the second category, the ideal observers [3–6], are optimal but limited only to search tasks with fixed display times. We are interested in ideal observers with unconstrained viewing times, which must solve optimally both the problem of *how* to accumulate evidence (for accuracy) and *when* to terminate the search (for speed).

One aspect that complicates the ideal observer analysis is that visual search tasks display great variability: the search could be *homogeneous*, where target and distractor appearances are known in advance and identical across locations, or *heterogeneous*, where some display properties are unknown and/or distinct across locations. Due to this variability, the corresponding ideal observers are often computationally intensive or even intractable. Despite the challenge we make progress in three directions. First, we propose a Bayesian framework that describes a lossless evidence accumulation procedure for homogeneous and heterogeneous search problems. Second, we describe a computational solution of ideal observers for visual search that is feasible for small set-sizes (number of search locations), and conjecture an efficient, analytical solution for arbitrary set-sizes. Last, we show that most current visual search models are sub-optimal. To our knowledge, our model is the first that can assess the optimality of

existing search models and humans. Moreover, our model characterizes the ER versus RT tradeoff as a function of main parameters that affect search difficulty (e.g. the distinctiveness of the target against the background clutter, the complexity of the image) as well as the degree of uncertainty in these parameters (e.g. varying image complexity from trial to trial).

Our work differs from previous ideal observer analyses in three aspects. First, we address the problem of finding the optimal time of decision, whereas most ideal observers based on signal detection theory [3–6] do not. Second, while ideal observers have been developed for discrimination [7, 8], search is inherently more difficult: input observations in visual search are high-dimensional, and lossless evidence accumulation occurs both locally at each display location and globally across the visual field. In addition, the state space for characterizing the optimal decision strategy is also high dimensional, which makes numerical solvers [9] for the ideal observer impractical and necessitates an analytical solution. Last, we do not model eye-movement, and focus instead on the simpler problem where the observer fixates at all times. This design choice reduces the dimensionality of the optimal decision strategy to a level where an exact solution is tractable. In comparison, models involving eye-movements [10, 11] suffer from the curse of dimensionality and thus are rarely optimal under our definition.

Results

Optimal evidence accumulation

Given sequential observations, an ideal observer is a model that achieves the best ER versus RT trade-off. The ideal observer consists of two components: a process to compute the posterior belief of “relevant” (discussed later) variables from the observations, and a process to decide when to stop making new observations. This section is about the first component. We start with existing theories on visual discrimination and homogeneous visual search, and extend them to account for general heterogeneous visual search.

Review of visual discrimination

We first review optimal evidence accumulation for visual discrimination. The stimulus is a single display item, either a distractor (denoted by $C = 0$ where C is the stimulus class) or a target ($C = 1$). The longer the stimulus is displayed, the more evidence an observer has to assess the target class. In this section we assume the display time is fixed, and the goal is to maximize accuracy given all available evidence. We assume that display items differ in a single attribute Y . For simplicity, throughout this paper we assume that the items are tilted bars, and the characteristic attribute is orientation (see Figure. 1(a); the model, however, is completely general and independent of the stimulus design). Target and distractor orientations can take values from two sets Θ_1 and Θ_0 , respectively. For example, $\Theta_1 = \{30^\circ, 10^\circ\}$, $\Theta_0 = \{20^\circ\}$ means that the distractor is 20° and the difference between target and distractor orientations (target contrast) is $\pm 10^\circ$.

The observer receives observation $X(t)$ that grows over time, and we can compute the log likelihood for each stimulus orientation $L_\theta(X(t)) \triangleq \log P(X(t)|Y = \theta)$. For example, if the input is a gaussian random walk, then $L_\theta(X(t))$ is a scaled version of the input (See Methods A.1 or [9]). Since here we assume that the optimal length of observation is given, we omit $X(t)$ ’s dependence on time and write it as X .

Bayesian inference computes the posterior belief of the stimulus class $P(C|X)$. It determines the theoretical error rate: e.g. if $P(C = 1|X) = 0.99$, then declaring target-present will produce a 1% error.

The posterior belief can be obtained deterministically from the *log likelihood ratio* $S \triangleq \log \frac{P(X|C=1)}{P(X|C=0)}$,

which is given by [7]:

$$S_{\text{discrim}}(X) = \mathcal{S}_{\theta \in \Theta_1}^{\max}(L_{\theta}(X) - \log(|\Theta_1|)) - \mathcal{S}_{\theta \in \Theta_0}^{\max}(L_{\theta}(X) - \log(|\Theta_0|)) \quad (1)$$

where $|\Theta_1|$ and $|\Theta_0|$ are the number of orientations in the target and the distractor set, respectively; $\mathcal{S}_{\max}(\cdot)$ is the “softmax” function, which is the marginalization operation in log probability space: it computes the log probability of a joint event from the log probabilities of its disjoint components:

$$\mathcal{S}_{\max}^{\mathcal{I}}(x) \triangleq \log \sum_{i \in \mathcal{I}} \exp(x_i)$$

Intuitively, since the target could take one of multiple mutually exclusive orientations, softmax combines evidence from each orientation $L_{\theta}(X)$ into the log probability for the target and for the distractor, respectively. The two log probabilities are then contrasted to yield the ratio $S_{\text{discrim}}(X)$.

Review of homogeneous visual search

In visual search, there are multiple (M) items on the display, and at most one of them can be a target. The task is to distinguish between target presence ($C = 1$) versus absence ($C = 0$). In homogeneous search both the target and distractor orientations $\Theta_1 = \{\theta_T\}$ and $\Theta_0 = \{\theta_D\}$ are distinct and unique. The log likelihood ratio is [5, 6]:

$$S_{\text{homo-search}}(X) = \mathcal{S}_{l=1 \dots M}^{\max}(L_{\theta_T}(X_l) - L_{\theta_D}(X_l)) - \log(M) \quad (2)$$

Since the target could appear at at most one location, softmax combines the local evidence at disjoint locations into a global log likelihood ratio $S_{\text{homo-search}}(X)$ (See [5] for derivation).

Search with unknown scene complexity or distractor type

In heterogeneous search, some search parameters (e.g. set-size, target contrast) are unknown before stimulus onset. Consequently, an ideal observer must infer these parameters on a trial-by-trial basis. We provide a general framework that encompasses the vast diversity within heterogeneous search, and a Bayesian inference procedure for lossless evidence accumulation.

To unify different heterogenous search tasks we need to define the *distractor mixture* ϕ . The distractor mixture encodes the distribution of stimulus orientation Y_l at any non-target location l : $\phi_i \triangleq P(Y_l = \theta_i | C_l = 0)$, $\forall l$. Many heterogeneous search tasks may be modeled by considering the effect of the search parameters on the distractor mixture (Fig. 1(e), examples below). As each search parameter setting affects the distractor mixture differently, ϕ is itself a random variable, with distribution $P(\phi)$, $\forall \phi \in \Phi$. We can then specify a heterogeneous search task using the distractor mixture distribution. For example:

- *Independent distractors* (Fig. 1(b)). Consider the case where the distractor at each location is sampled independently and uniformly at random from a set of n possible values [6] (e.g. looking for a key on a desk full of different objects). Since at any non-target location, the orientation Y_l could take any of the n values with probability $P(Y_l = \theta | C_l = 0) = 1/n$, $\forall l$, this setting can be described using only one distractor mixture $\Phi = \{\phi\}$ where $\phi : \phi_i = \frac{1}{n}, \forall i$. (the corresponding inference [6] is thus subsumed by our analysis).
- *Tied distractors* (Fig. 1(c)). Consider the case where the distractors are constrained to be *identical* across locations (e.g. looking for a camouflaged prey in a field of tall grass with uniform texture). Here the distractor orientations are deterministic only given the common orientation, which is randomly chosen from n possibilities. More specifically, the distractor mixture set uses n distractor mixtures: $\Phi = \{\phi^{(k)}\}_{k=1}^n$, where each distractor mixture specifies a deterministic distribution $\phi^{(k)} : \phi_i^{(k)} = \mathbb{I}(k = i)$. For example, if $n = 3$, then $\Phi = \{[1 \ 0 \ 0], [0 \ 1 \ 0], [0 \ 0 \ 1]\}$.

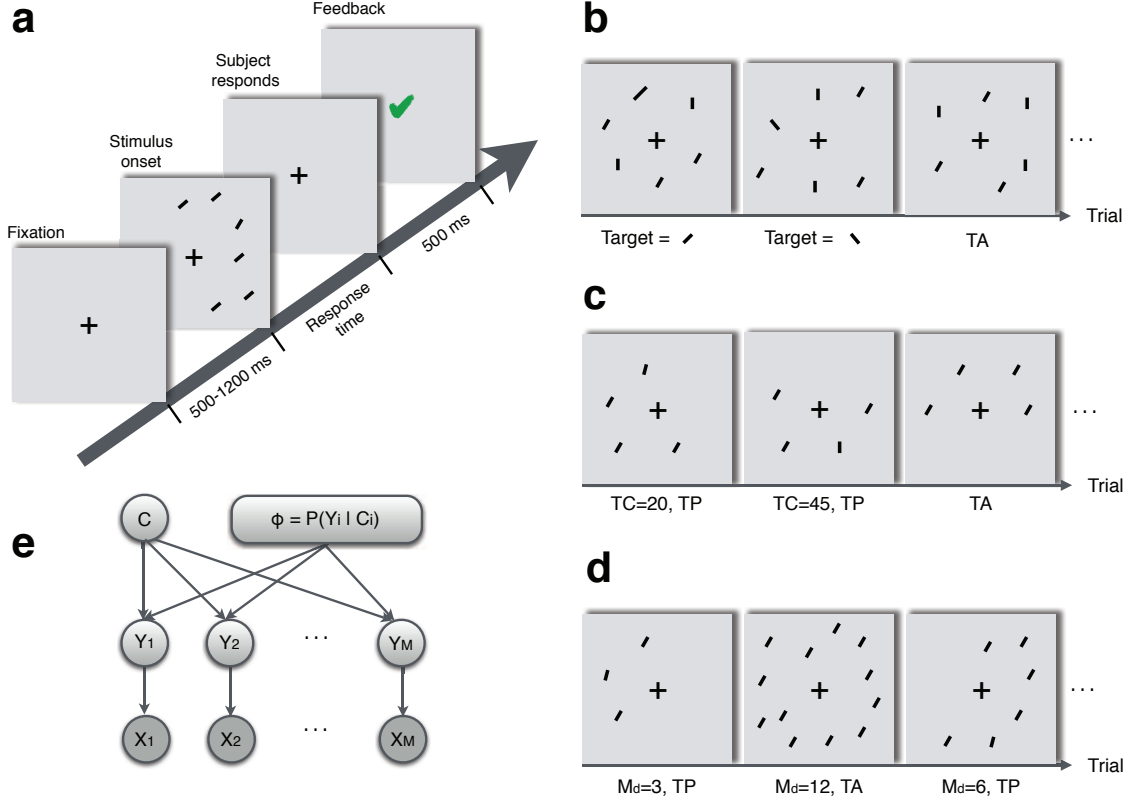


Figure 1. Heterogeneous visual search. (a) Experiment setup. Each trial starts with a fixation screen. Next, a ‘stimulus’ is displayed. The stimulus is an image containing M oriented bars, one of which may be the target. The stimulus disappears as soon as the subject responds by pressing one of two keys, to indicate whether a target is detected or not. Then the screen shows feedback of whether the response is correct, which concludes the trial. Fixation at the center is enforced at all times. (b-d) Various heterogeneous search designs. (b) Target and distractor orientations change independently from trial to trial; (c) distractor orientation is tied across all locations, but the target-distractor orientation contrast (TC) varies; (d) mixed number of display items M_d . TP: target present; TA: target absent. (e) A general graphical model (see details in Equation. 3,4,5 and Methods. A.2) that explains different designs in (b-d).

- *Unknown scene complexity* (Fig. 1(d)). In the case display items could appear at M_d randomly selected locations out of a maximum of M , where M_d , the set-size, is sampled uniformly from a discrete set of n values. In this case we introduce “blanks” (denoted θ_{\emptyset}) as a special type of distractor. Any display pattern can be treated as M_d regular items and $M - M_d$ blanks¹. Thus we

¹In this case the inference procedure is biased: the actual distribution of display locations consists of delta functions at the possible M_d values, while the proposed model implies a mixture of Multinomial distributions, whose means match one-to-one the possible M_d values. Thus the implied distribution is a fuzzier version of the real world, leading to conservative estimates of the log likelihood ratio. In practice, however, blank differentiation is typically instantaneous (i.e. with large diffusion slope), so the consequence of the bias is negligible.

need n distractor mixtures, each encodes a distractor probability proportional to the corresponding set-size: $P(Y_l \in \Theta_0 | C_l = 0) = M_d/M$. For example, say $M = 7$, and M_d could be either 1 or 3, then $\Phi = \{[1/7, 6/7], [3/7, 4/7]\}$.

Given the distribution of distractor mixtures $P(\phi)$, the log likelihood ratio $S_{\text{hetero-search}}$ is:

$$S_{\text{hetero-search}}(X) = \mathcal{S}_{l=1 \dots M}^{\max} (L_{\Theta_1}(X_l) - L_{\Theta_0|X}(X_l)) - \log(M) \quad (3)$$

$$\text{where } L_{\Theta_1}(X_l) = \mathcal{S}_{\theta \in \Theta_1}^{\max} (L_{\theta}(X_l)) - \log(|\Theta_1|) \quad (4)$$

$$L_{\Theta_0|X}(X_l) = -\mathcal{S}_{\phi \in \Phi}^{\max} \left(-\mathcal{S}_{\theta \in \Theta_0}^{\max} (L_{\theta}(X_l) + \log \phi_{\theta}) + Q_{\phi}(X) \right) \quad (5)$$

and $Q_{\phi}(X) \triangleq \log P(\phi|X)$ is the posterior belief of distractor mixture ϕ given observations from all locations (details later).

The log likelihood ratio expression above is obtained by nesting properly equations from before. At the highest level $S_{\text{hetero-search}}$ (Eq. 3) is reminiscent of $S_{\text{homo-search}}$ (Eq. 2). In addition, since target orientation is unknown, it must be inferred in Equation. 4 (implemented via a soft-max, as in Equation. 1). The same applies to distractor orientation, except that the uncertainty is two-fold: both the value and the distribution of distractor orientation are unknown. Hence, two soft-maxes (Eq. 5), one over the distractor orientation Y_l and the other over the distractor mixture ϕ , are necessary.

The distractor mixture ϕ can be inferred using gain-control as follows:

$$Q_{\phi}(X) = \text{gain-control} \left(\log P(\phi) + \sum_{l=1}^M \mathcal{S}_{\theta \in \Theta_0}^{\max} (L_{\theta}(X_l) + \log \phi_{\theta}) \right)$$

$$\text{gain-control}(A_i) = A_i - \mathcal{S}_j^{\max} (A_j)$$

In conclusion, heterogeneous visual search is challenging because neither the distractor orientation nor its distribution is known in advance. Fortunately, they can be inferred from evidence over the entire visual field, after which the problem reduces to the well-understood, homogeneous visual search. The proposed inference framework incorporates different heterogeneous search modalities such as mixed target type, mixed distractor type and unknown scene complexity.

Optimal decision strategy

While the log likelihood makes the best use of existing observations to minimize ER, the ideal observer also requires a stopping strategy in order to trade off RT with ER. The stopping criterion is context-sensitive: an ideal observer should discern whether a situation emphasizes urgency (e.g. competing with other predators for prey) or caution (e.g. avoiding poisonous mushrooms). One common way to capture the importance of error versus time is through a single risk function:

$$\text{Risk} = ER + C_{\text{time}} \text{average RT} = \sum_{C=0}^1 (P_C(D \neq C) + C_{\text{time}} \mathbb{E}_C[T]) \quad (6)$$

where $D \in \{0, 1\}$ is the observer's decision and $T \in [0, \infty]$ is the response time, respectively, of the sequential test. C_{time} is the relative cost of time with respect to error². The optimal test achieves the lowest risk among all tests.

²For simplicity we assume that false positives and false negatives have the same cost, and so do the response times under each class. Different costs can be easily accommodated without affecting the overall analysis.

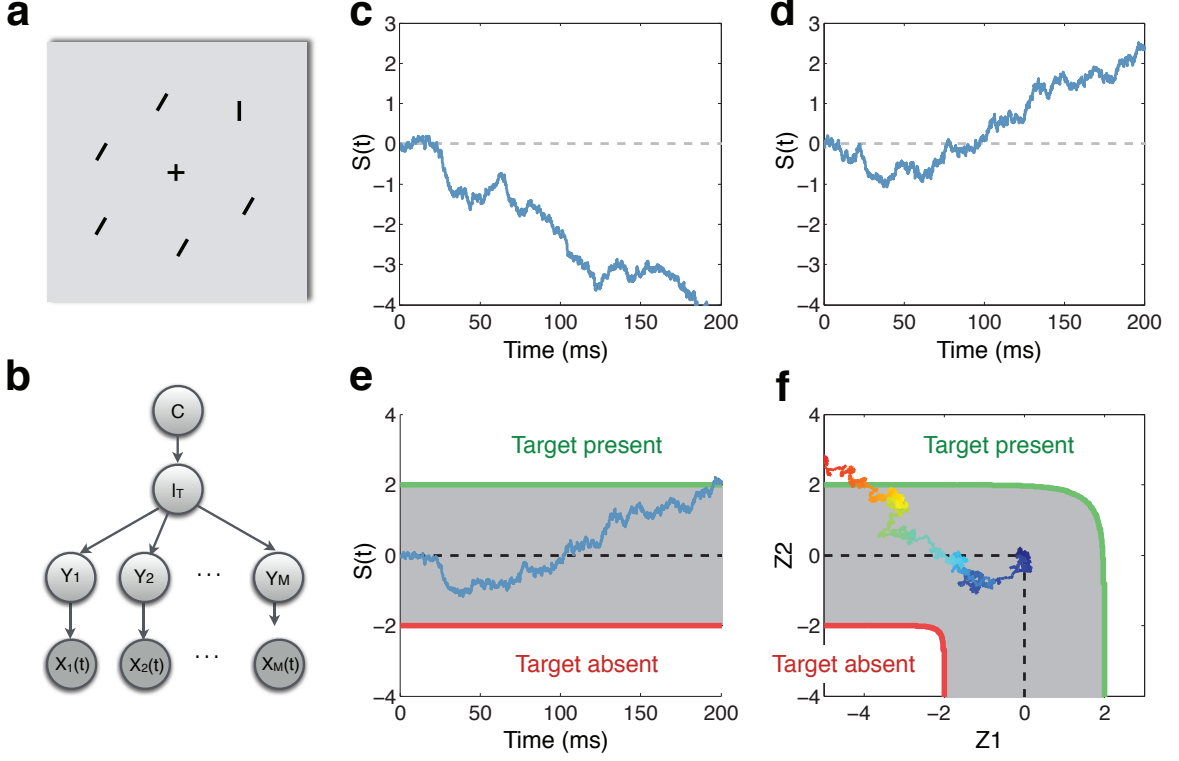


Figure 2. Decision strategies for homogeneous visual search. (a) A homogeneous visual search task: spot the 0° target amongst 20° distractors. (b) A generative model of the problem, simplified from Figure 1(e). The class variable C (target present or not) controls the location l_T of the target ($l_T = 0$ means target absent). l_T in turn determines the item type Y_l at each location. Given the item type, the observations X_l are independent in space and i.i.d. in time. To perform probabilistic inference, each location computes the local log likelihood ratio $Z_l = \log \frac{P(X_l|C_l=1)}{P(X_l|C_l=0)}$ over time: (c) Z_l at a distractor location, (d) Z_l at the target location. Two decision strategies that make use of the probabilistic interpretation for a two-dimensional visual search problem. SPRT (e) thresholds the one-dimensional log likelihood ratio $S(t)$ (Eq. 2), whereas the optimal (f) uses a decision boundary in the joint space of $\{Z_1, Z_2\}$. Time in (e-f) is color-coded, cooler colors means earlier.

Two components are necessary to describe the optimal test: a state space $\vec{Z}(t)$ over time and a decision strategy that associates each state and time with an action. One common constraint on the state space is that it must be Markov in time: $\vec{Z}(t)$ must be sufficient in summarizing past observations so that given $\vec{Z}(t)$, future observations become independent from the past (see Methods A.3). Once this constraint is satisfied, the problem may be formulated as a partial observation Markov decision process (POMDP) [13], and the optimal strategy may be solved exactly using dynamic programming.

We choose $\vec{Z}(t)$ to be the collection of log likelihood ratios from all locations:

$$\vec{Z}(t) : Z_l(t) = L_{\theta_T}(X_l) - L_{\theta_D}(X_l), \quad l = 1 \dots M$$

The decision strategy depends on the characteristics of the input X . We consider the most common

formulation of input as a gaussian random walk at each location (e.g. [9, 12]). The input is parameterized by the *drift-rate* $\mu_{C,l}$, which depends on the stimulus class C and the location l (Fig. 2(c-d)). A larger drift-rate difference between the two classes $|\mu_{1,l} - \mu_{0,l}|$ implies a higher signal-to-noise ratio, or equivalently, an easier discrimination problem at location l .

Computational solution for low-dimensional problems

The optimal decision may be computed numerically using dynamic programming [13, 14]. Define $R(\vec{Z}(t), t)$ as the lowest total risk an observer could incur starting from $\vec{Z}(t)$ at time t . The optimal risk is equivalent to $R(\vec{0}, 0)$, the total risk from time 0 onwards with a flat prior. $R(\vec{Z}, t)$ is recursively given by:

$$R(\vec{Z}(t), t) = \min \begin{cases} 1 - P_0(\vec{Z}(t)) & D = 0: \text{declare target absent} \\ P_0(\vec{Z}(t)) & D = 1: \text{declare target present} \\ C_{\text{time}}\delta t + \mathbb{E}_{\vec{Z}(t+\delta t)|\vec{Z}(t)} R(\vec{Z}(t+\delta t), t+\delta t) & D = \emptyset: \text{wait} \end{cases} \quad (7)$$

At any time t and any state $\vec{Z}(t)$, the ideal observer picks the action $D \in \{\emptyset, 0, 1\}$ that yields the lowest risk. If declaring target-absent, the observer makes a false rejection mistake. The false reject probability can be computed from the state $\vec{Z}(t)$ and is denoted $P_0(\vec{Z}(t))$ (see Methods A.5 and Equation 14). If waiting for more evidence, the observer trades off the cost $C_{\text{time}}\delta t$ for a new observation of duration δt , and access to the cumulative risk at $t + \delta t$.

The optimal decision strategy is defined over a $M + 1$ dimensional state-space. The state space is separated by decision boundaries/surfaces into three different decision regions [15]. Furthermore, the recurrence equation 7 is time invariant. As a result, the optimal decision is constant in time (see [13] and Methods A.5) and the decision surfaces have $M - 1$ dimensions. This showcases the difference between the optimal decision strategy for visual search and that for visual discrimination [7, 9]. In discrimination, knowing the target-present probability given past observation is sufficient to compute the likelihood of new observations, hence the state-space is *always one-dimensional*. On the contrary, visual search requires the additional knowledge of how the target-present probability breaks down to each location, hence the dimensionality of the state-space is proportional to the set-size.

Conjecture for high-dimensional problems

In homogeneous discrimination, the optimal decision strategy is given by the classical Sequential Probability Ratio Test (SPRT) [16], which compares the instantaneous log likelihood ratio $S(X(t))$ to a pair of constant thresholds. This strategy has been proven only asymptotically optimal [17, 18] for problems that involve multiple hypotheses, such as visual search. Nonetheless, in the case of homogeneous visual search, we conjecture that the optimal decision strategy for high-dimensional search is similar to SPRT: in fact it uses two SPRTs defined on scaled log likelihood ratios.

Conjecture 1. (*Uniform drift-rates*) If all locations share the same drift-rate ($\mu_{1,l} = -\mu_{0,l} = \mu, \forall l$), let τ_+ and τ_- be the optimal upper and lower thresholds for visual discrimination at location l with a cost of time of C_{time} , then the optimal decision surfaces for minimizing the risk function in Equation. 6 with the same cost of time C_{time} are:

$$S_+(\vec{Z}(t)) = \frac{1}{a_+} \mathcal{S}_{\max_{l=1, \dots, M}} (a_+(Z_l(t) - \log(M))) \geq \tau_+ \quad (8)$$

$$S_-(\vec{Z}(t)) = \frac{1}{a_-} \mathcal{S}_{\max_{l=1, \dots, M}} (a_-(Z_l(t) - \log(M))) \leq \tau_- \quad (9)$$

where a_+ and a_- are unknown parameters.

Conjecture 1 states that the optimal decision strategy is to wait until either $S_+(X(t)) \geq \tau_+$ ($D = 1$) or $S_-(X(t)) \leq \tau_-$ ($D = 0$). The thresholds τ_+ and τ_- are obtained easily by solving a one-dimensional dynamic programming problem [9]. The thresholds are chosen to guarantee asymptotic optimality. Intuitively, when there is only one location ($M = 1$), Conjecture. 1 reduces to SPRT, which is optimal for visual discrimination. More trickier is the asymptotic case where $M > 1$ (multiple locations) but the decision is effectively reduced to concerning only one location l^* . This is when all $Z_l(t)$'s at locations $l \neq l^*$ become arbitrarily close to $-\infty$. i.e. when the observer has accumulated a significant amount of information to rule out location $l \neq l^*$ as the target location. In this case:

$$S_+(\vec{Z}(t)) = \frac{1}{a_+} \mathcal{S}_{\max_{l=1, \dots, M}} (a_+(Z_l(t) - \log(M))) \approx Z_{l^*}(t) - \log(M)$$

$$S_-(\vec{Z}(t)) \approx Z_{l^*}(t) - \log(M)$$

This case happens at location l^* with a probability of $1/M$ when the target is present, and a probability of 1 when absent. Hence, asymptotically the visual search problem reduces to a visual discrimination problem at location l^* with a log prior ratio of $\log(1/M)$. The best ER vs RT trade-off is achieved when $Z_{l^*}(t)$ is compared against thresholds adjusted by the log prior ratio: $\tau_+ + \log(M)$ and $\tau_- + \log(M)$ (for proof see Methods A.4), which is exactly the asymptotic behavior of Equation. 8 and 9.

Figure. 3(a-b) and Figure. 4 show excellent empirical match between the conjectured thresholds and the optimal thresholds in 2D.

Our conjecture can be extended to handle a particular heterogeneous search task where the display properties are known, but the drift-rates are different across locations. We refer to this as *heterogeneous drift-rate search*, to distinguish it from the much more complex heterogeneous search discussed before. For heterogeneous drift-rate search, Conjecture. 1 can be extended by introducing a correction factor for each location:

Extension 2. (*Non-uniform drift-rates*) Let $\tau_+^{(l)}$ and $\tau_-^{(l)}$ be the upper and lower thresholds for visual discrimination with a time cost of C_{time} at location l , define $c_+^{(l)} = \tau_+^{(M)}/\tau_+^{(l)}$ and $c_-^{(l)} = \tau_-^{(M)}/\tau_-^{(l)}$, the optimal decision surface for visual search with the same time cost is:

$$S_+(\vec{Z}(t)) = \frac{1}{a_+} \mathcal{S}_{\max_{l=1, \dots, M}} \left(c_+^{(l)} a_+(Z_l(t) - \log(M)) \right) \geq \tau_+^{(M)} \quad (10)$$

$$S_-(\vec{Z}(t)) = \frac{1}{a_-} \mathcal{S}_{\max_{l=1, \dots, M}} \left(c_-^{(l)} a_-(Z_l(t) - \log(M)) \right) \geq \tau_-^{(M)} \quad (11)$$

The only difference from the uniform drift-rate case (Eq. 10) is that we scale the local diffusions by a location-dependent factor $c_+^{(l)}$. This factor normalizes the diffusion at each location by its relative efficiency with respect to a reference location (arbitrarily chosen to be location M). Similar to the intuition that justifies Conjecture. 1, in the asymptotic case where only one location l^* is relevant, $S_+(\vec{Z}(t)) \approx \tau_+^{(1)}(Z_{l^*}(t) - \log(M))/\tau_+^{(l^*)}$. Therefore when the optimal discrimination ER vs RT tradeoff is reached at $Z_{l^*}(t) - \log(M) = \tau_+^{(l^*)}$, the visual search log likelihood ratio $S_+(\vec{Z}(t))$ is at the prescribed threshold $\tau_+^{(M)}$ (Eq. 10). The same argument applies to the lower threshold $\tau_-^{(l^*)}$ and Equation. 11.

Extension. 2 only requires solving M one-dimensional dynamic programming problems for $\tau_+^{(l)}$ and $\tau_-^{(l)}$, which is more scalable than the optimal procedure (Eq. 7) that scales exponentially with M . Figure. 3(c-d) shows that the predicted thresholds from Conjecture. 2 match the optimal thresholds from dynamic programming in 2D for a variety of costs of time and drift-rates.

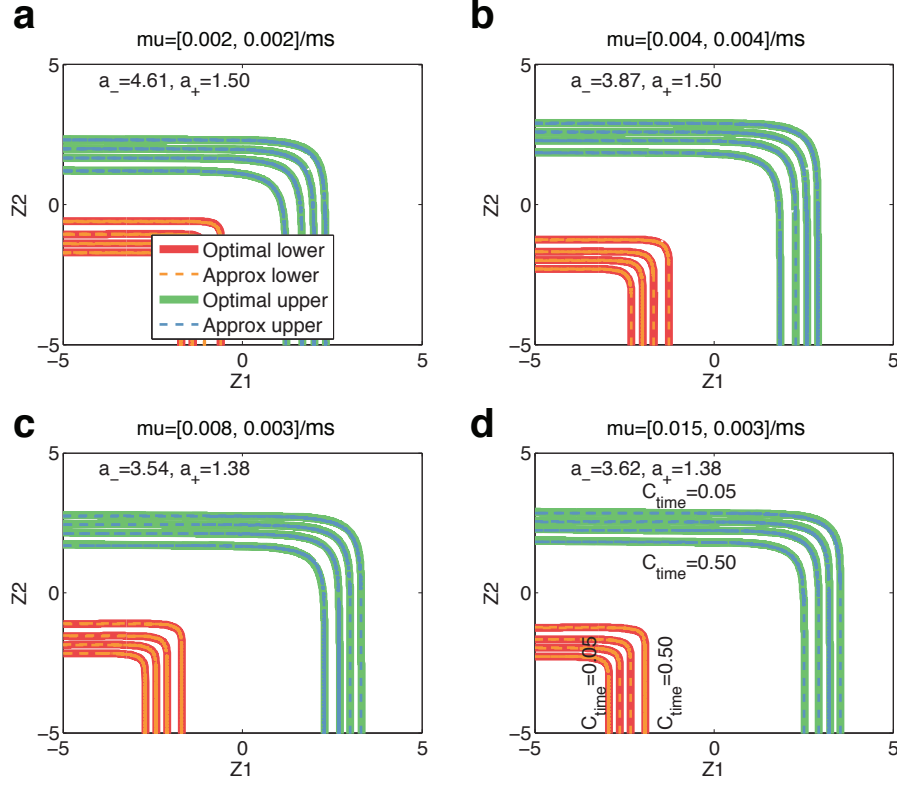


Figure 3. Optimal sequential test for 2D visual search. (a-b) Optimal decision thresholds and approximations for different costs of time $C_{\text{time}} \in \{0.5, 0.2, 0.1, 0.05\}$ in homogeneous search. Decision boundaries are approximated using Equation. 10 and 11 with $a_+ = 1.50$ and $a_- = 4.61$. (c-d) Optimal decision thresholds and approximations for heterogeneous drift-rate search. Drift-rates are (a-b) $\pm 2/\text{sec}$, (c) $\{\pm 8, \pm 3\}/\text{ms}$ and (d) $\{\pm 15, \pm 3\}/\text{sec}$.

Optimality analysis of current search models

How are existing visual search models compare against the ideal observer? For fairness we compare only approaches that perform probabilistic inference on the graphical model in Figure. 2 (b). These approaches, listed below, differ only in the decision strategy [3]:

- **a-SPRT** (Fig. 2(f)): our two-SPRT approach that uses two decision surfaces prescribed in Conjecture. 1 and Extension 2 to approximate the ideal observer.
- **SPRT** [5] (Fig. 2(e)): a Bayesian extension of Ward’s SPRT [16] into testing composite hypotheses. SPRT compares the log likelihood ratio of target-present versus target-absent $S(t)$ (Eq. 2) against a pair of thresholds. Since the SPRT is subject to the same asymptotic analysis in Conjecture. 1, it uses the same thresholds τ_- and τ_+ as does the a-SPRT. Essentially, SPRT is a special case of Equation. 8 and 9 where $a_+ = a_- = 1$.
- **SPRT-opt**: the same as SPRT above except that it optimizes the upper and lower thresholds to minimize the risk function (Eq. 6). Since SPRT-opt may use different thresholds from those in

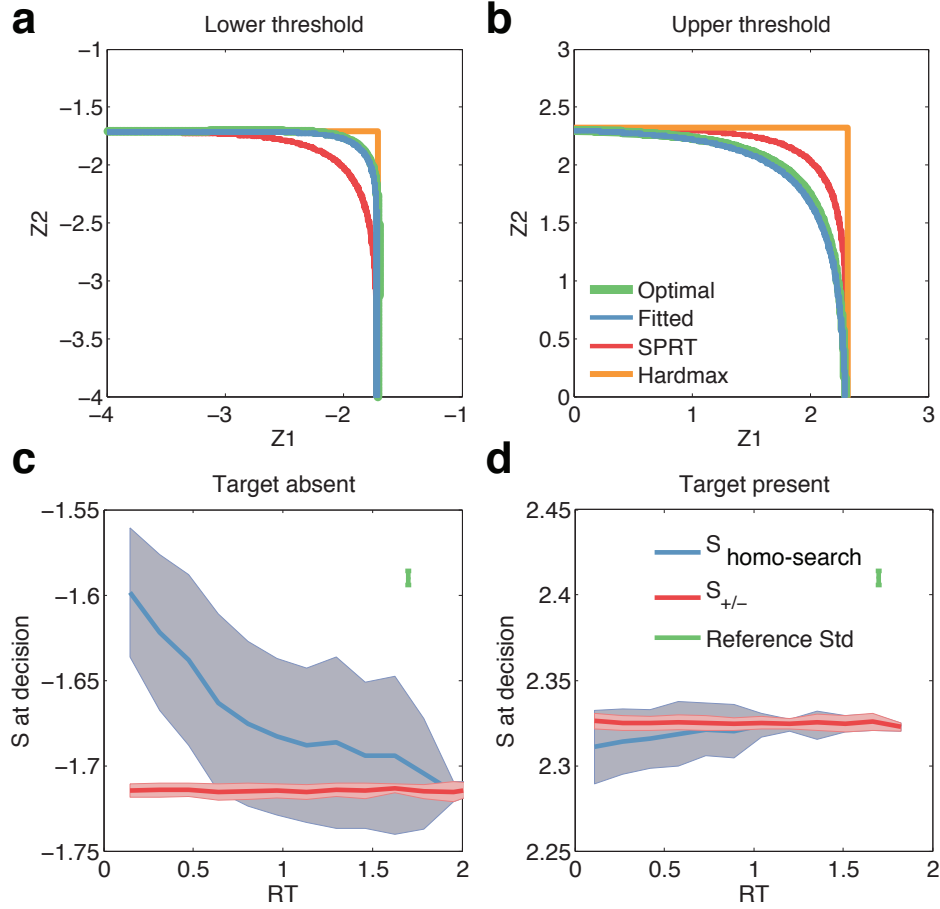


Figure 4. Sequential testing strategies for homogeneous visual search in two-dimensions. The optimal and various alternative decision strategies are compared in terms of (a) the lower and (b) the upper threshold in the joint space of $\{Z_1, Z_2\}$. The a-SPRT thresholds are obtained from Equation. 8 and 9 with $a_+ = 1.5$ and $a_- = 3.9$; Both SPRT and Hardmax use the optimal threshold for visual discrimination so that asymptotically they are consistent with the optimal strategy. Input to each display location has a drift-rate of $\pm 4/sec$. (c-d) Each panel shows the log likelihood ratio $S(t)$ distribution at the time of decision under the optimal decision strategy from $1k$ Monte-Carlo simulations. As references, the distribution of S_- when target is absent (c) and of S_+ when present (d) are shown. S_{\pm} is not deterministic because time is discretized in the simulation, which causes the log likelihood ratios to have finite-sized jumps. Standard deviation of the jumps are shown as another reference. Drift-rate of the observation is $\pm 2/sec$.

the regular SPRT, it may not be asymptotically optimal. However, this does not prevent SPRT-opt from outperforming the regular SPRT (which is asymptotically optimal). This is because the asymptotic (i.e. long) decisions may only take up a tiny fraction of all the decisions (especially in easy tasks), and SPRT-opt may do better by focusing on the risk for shorter decisions.

- **Hardmax** [5, 11]: an efficient approximation to SPRT. Each location decides whether it contains a target ($D_l = 1$) or a distractor ($D_l = 0$) based solely on the local belief $Z_l(t)$. The observer declares target-present when any location reports a target detection, declares target-absent when

all locations report a distractor, and waits for more information otherwise. Hardmax is also a special case of Equation. 8 and 9 where $a_+ = a_- = \infty$.

Decision surfaces comparison. We want to see how these approaches differ from the optimal in various aspects. First, how different are their decision surfaces? In Figure. 4(a-b), we compare them on a visual search task with two display locations where it is computationally feasible to solve for the optimal decision boundary using dynamic programming. Since the decision boundaries are constant in time, they can be visualized in the 2-D space of Z_1 and Z_2 only. Each decision boundary is of the form $\{(Z_1, Z_2) | S(Z_1, Z_2) = \tau\}$, i.e. all pairs of Z_1 and Z_2 that could make the log likelihood ratio S reach a threshold of τ .

We observe that both the Hardmax and SPRT *differ significantly* from the optimal in terms of the decision surfaces (Fig. 4(a-b)). SPRT is conservative, because both thresholds bend outwards with respect to the optimal thresholds, which translates to longer decision times for both target-present and target-absent runs. Hardmax, on the other hand, is faster in declaring target-absent but slower in declaring target-present.

Can time-varying threshold make SPRT optimal? A common practice in modeling decision making in visual discrimination is to employ a time-varying threshold. Can the optimal decision mechanism for visual search also be implemented using SPRT-opt with a *time-varying* threshold? We reject this hypothesis by computing the $S(t)$ distribution at the time of decision under the optimal test (Fig. 4(c-d)). If a time-varying threshold exists on $S(t)$ to recover the optimal strategy, the $S(t)$ values should be unique at the time of decision. Instead, we observe a wide spread in the $S(t)$ distribution. Therefore, $S(t)$ is not a sufficient statistic to implement the optimal test, and SPRT is sub-optimal in visual search [16].

Risk comparison. The decision surfaces comparison above has one caveat: we consider all places on the decision boundary where decisions *could* be taken, ignoring the fact that some places on the boundary are more likely to be reached than others in an actual decision task. E.g. consider Figure. 4b when the search task is easy, the diffusions when the target is present will most likely fall in the region of $\{Z_2 > 0, Z_1 \ll 0\}$ and $\{Z_2 \ll 0, Z_1 > 0\}$, and rarely visit the region of $\{Z_1 > 0, Z_2 > 0\}$ where the difference among the strategies is the most noticeable. This reasoning suggests that we should compare these strategies in terms of their actual *risk* value.

The risks for the strategies in a homogeneous search task are shown in Figure. 5. Hardmax and SPRT are highly sub-optimal. SPRT-opt is almost indistinguishable from a-SPRT in the low time-cost scenario, but becomes sub-optimal when the cost of time becomes very high, i.e. when the decision time is short. Although we have not yet proven that a-SPRT is optimal, it is sufficient to conclude that any model that underperforms it is sub-optimal.

In heterogeneous drift-rate search (Fig. 6), we see that even with two display locations, both SPRT-opt and Hardmax³ are suboptimal when the drift-rates differ significantly across locations. The sub-optimality becomes progressively more pronounced as the heterogeneity of drift-rates increases. Behaviorally, when the drift-rate heterogeneity is large, Hardmax achieves near-identical ER vs RT trade-offs at both locations, whereas SPRT-opt and a-SPRT learn to sacrifice the ER at the low drift-rate location for a faster RT overall (Fig. 6c).

In conclusion, decision strategies employed by existing search models are sub-optimal. Hardmax, where one combines local decisions to reach a global decision, is sub-optimal in almost all scenarios. The SPRT-opt, where one executes a one-dimensional SPRT with optimized thresholds, is near-optimal in low cost, homogeneous search scenarios. When the cost of time is high and when the drift-rate is heterogeneous across locations, the SPRT-opt becomes sub-optimal, but remains similar to the optimal strategy in terms of ER and RT.

³We do not include SPRT because it is not clear how to condense the M asymptotically optimal thresholds, one for each decision surface, into just one for the SPRT. Instead we trust that SPRT-opt, with the ability to optimize the thresholds, should always outperform any SPRT

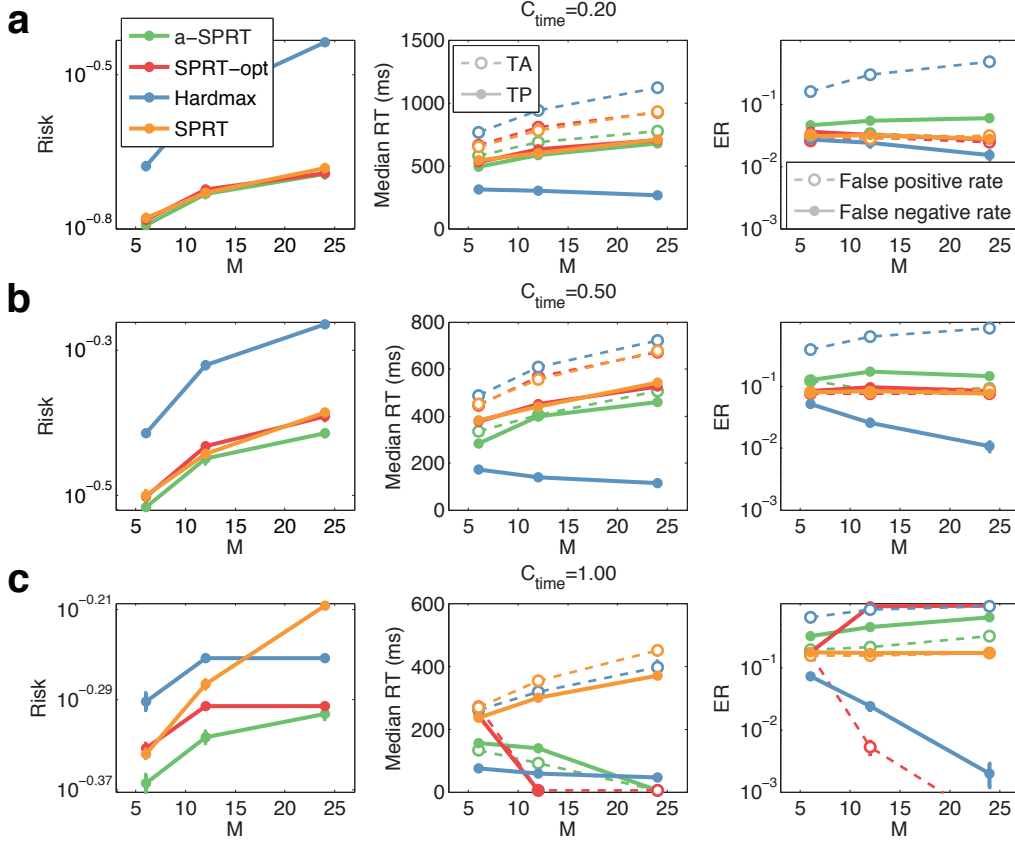


Figure 5. Risk comparison of common decision strategies in homogeneous visual search. a-SPRT, SPRT-opt, SPRT and Hardmax are compared under different costs of times (a) $C_{\text{time}} = 0.2$, (b) $C_{\text{time}} = 0.5$, and (c) $C_{\text{time}} = 1$ with a drift-rate of $\pm 12/\text{sec}$. Hardmax is sub-optimal in all cases. Regular SPRT is sub-optimal in the high cost scenario. SPRT-opt slightly under-performs a-SPRT in terms of the risk. a-SPRT and SPRT-opt are similar in terms of the RT during target-present (TP) and target-absent (TA), as well as the false positive rate and the false negative rate. Error bars are one standard error computed from $10k$ runs.

Discussion

We have described the ideal observer model to optimally trade off RT and ER in visual search. The ideal observer relies on lossless evidence accumulation and optimal decision making. We accumulate evidence according to Bayesian inference, where unknown search parameters such as set-size, target contrast and target location are marginalized out. Our model can handle a variety of heterogeneous visual search tasks by modeling the distractor mixture, which is the distribution of distractor orientations. Our model augments the basic search network proposed in [5] with a parallel, gain-control circuit that specializes in marginalizing out the distractor mixtures. While the basic search network is common to all tasks, the gain-control circuit is task dependent. This functional separation provides robustness: our model can simply adjust the distractor mixture distribution $P(\phi)$ to switch between different search tasks.

We have conjectured a novel procedure, a-SPRT, to compute the optimal decision policy for high-dimensional homogeneous visual search and heterogeneous drift-rate search. The a-SPRT makes use of

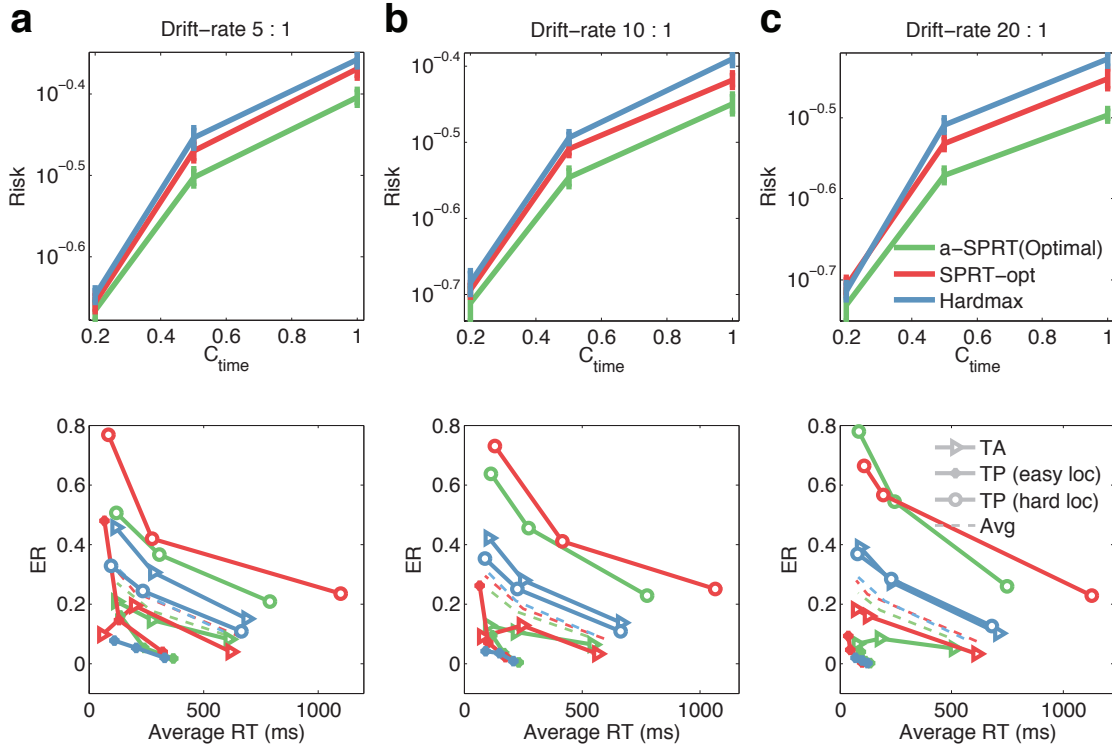


Figure 6. Risk comparison of common decision strategies in heterogeneous drift-rate visual search. a-SPRT, SPRT-opt and Hardmax are compared under various costs of time. The first row shows the overall risk versus the cost of time. The second row shows the ER vs RT tradeoff under different costs of time (dots) and under three separate conditions (lines): target-absent (TA), target-present (TP) at the location with a larger drift-rate (easy) and target-present at the hard location. Drift-rates are **(a)** $\{\pm 5, \pm 1\}/\text{sec}$, **(b)** $\{\pm 10, \pm 1\}/\text{sec}$ and **(c)** $\{\pm 20, \pm 1\}/\text{sec}$. One standard error in both RT and ER computed from 1k runs are shown but too small to be visible. Both SPRT-opt and Hardmax underperform the optimal test.

two one-dimensional SPRTs with different scaling factors, and with thresholds that are constant in time. In two dimensions, the resultant decision boundary matches closely that of the optimal strategy. The conjecture is preferable over the standard dynamic programming procedure, which does not scale to high (more than three) dimensions.

Our model enables optimality analysis of humans and phenomenological models of visual search in certain conditions. We compared various plausible decision strategies and discover that most of them are sub-optimal. While the SPRT with optimized thresholds behaves similarly as the proposed strategy in homogeneous search tasks, it is sub-optimal in search tasks where the signal-to-noise ratio is heterogeneous across locations.

We highlight several unsolved issues for future work. First, it remains an open question why the optimal decision boundaries for homogeneous search can be described by two scaled-SPRTs. Second, we do not know how the scaling factors a_+ and a_- depend on search parameters, and therefore must

search numerically for their values to minimize the risk. A better understanding is required to generalize ideal observers of visual search into greater dimensionality and heterogeneity. Third, since both SPRT-opt and a-SPRT can be efficiently implemented using neural hardware [5, 19], if Conjecture. 1 is proven true, it would imply that the computations required by the ideal observer *can* be carried out in cortex. Last, in light of the marked difference between alternative models and the ideal observer in the case of heterogeneous drift-rate search, it would be interesting to test subjects in this case to see which model best captures human behavior, and whether humans *are* indeed optimal.

References

1. Ratcliff R (1985) Theoretical interpretations of the speed and accuracy of positive and negative responses. *Psychological Review* 92: 212.
2. Purcell BA, Schall JD, Logan GD, Palmeri TJ (2012) From salience to saccades: multiple-alternative gated stochastic accumulator model of visual search. *The Journal of Neuroscience* 32: 3433–3446.
3. Palmer J, Verghese P, Pavel M (2000) The psychophysics of visual search. *Vision Research* 40: 1227–1268.
4. Verghese P (2001) Visual search and attention: A signal detection theory approach. *Neuron* 31: 523–535.
5. Chen B, Navalpakkam V, Perona P (2011) Predicting response time and error rates in visual search. In: *NIPS*. pp. 2699–2701.
6. Ma WJ, Navalpakkam V, Beck JM, Van Den Berg R, Pouget A (2011) Behavior and neural basis of near-optimal visual search. *Nature neuroscience* 14: 783–790.
7. Palmer J, Huk A, Shadlen M (2005) The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision* 5.
8. Gold JJ, Shadlen MN (2001) Neural computations that underlie decisions about sensory stimuli. *Trends in cognitive sciences* 5: 10–16.
9. Drugowitsch J, Moreno-Bote R, Churchland AK, Shadlen MN, Pouget A (2012) The cost of accumulating evidence in perceptual decision making. *The Journal of Neuroscience* 32.
10. Chukoskie L, Snider J, Mozer MC, Krauzlis RJ, Sejnowski TJ (2013) Learning where to look for a hidden target. *Proceedings of the National Academy of Sciences* 110: 10438–10445.
11. Najemnik J, Geisler WS (2008) Eye movement statistics in humans are consistent with an optimal search strategy. *Journal of Vision* 8: 4.
12. Thornton TL, Gilden DL (2007) Parallel and serial processes in visual search. *Psychological review* 114: 71.
13. Cassandra AR, Kaelbling LP, Littman ML (1994) Acting optimally in partially observable stochastic domains. In: *AAAI*. volume 94, pp. 1023–1028.
14. Bellman R (1956) Dynamic programming and lagrange multipliers. *Proceedings of the National Academy of Sciences of the United States of America* 42: 767.

15. Sobel M, et al. (1953) An essentially complete class of decision functions for certain standard sequential problems. The Annals of Mathematical Statistics 24.
16. Wald A (1945) Sequential tests of statistical hypotheses. The Annals of Mathematical Statistics 16: 117–186.
17. Dragalin VP, Novikov AA (1987) Asymptotic solution of the kiefer-weiss problem for processes with independent increments. Theory of Probability & Its Applications 32.
18. Schwarz G, et al. (1962) Asymptotic shapes of bayes sequential testing regions. The Annals of mathematical statistics 33.
19. Jazayeri M, Movshon JA (2006) Optimal representation of sensory information by neural populations. Nature neuroscience 9.

A Methods

A.1 Gaussian input

We assume that at each location, the visual system receives gaussian instantaneous observation $\dot{x}(t)$. The cumulative observation at location l : $X_l(t) = x(t)$ is a Gaussian random walk:

$$\begin{aligned}\dot{x}(t) &\sim \mathcal{N}(\mu_\theta \delta t, \delta t) \\ x(t + \delta t) &= x(t) + \dot{x}(t) \delta t\end{aligned}$$

where μ_θ , the slope of the random walk, depends on the stimulus orientation θ .

At any location l , both the log likelihood $\log P(X_l|Y_l = \theta)$ and the log likelihood ratio $\log \frac{P(X_l|Y_l = \Theta_1)}{P(X_l|Y_l = \Theta_0)}$ are linear functions of $x(t)$. First observe that the evidence X at time $t = K\delta t$ is distributed as a gaussian:

$$\begin{aligned}x(t = K\delta t) &= \sum_{k=1}^K \dot{x}(k\delta t) \delta t \sim \mathcal{N}(\mu_\theta t, t) \\ L_\theta(x(t)) &= \log P(x(t)|Y_l = \theta) = \log \mathcal{N}(x(t)|\mu_\theta t, t) \\ &= \mu_\theta x(t) + \frac{\mu_\theta^2 t}{2} + \text{Const}\end{aligned}\tag{12}$$

where Const is independent of θ . Similarly,

$$\begin{aligned}Z_l(t) &= \log \frac{P(x(t)|Y_l = \Theta_1)}{P(x(t)|Y_l = \Theta_0)} \\ &= L_{\Theta_1}(x(t)) - L_{\Theta_0}(x(t)) \\ &= (\mu_{\Theta_1} - \mu_{\Theta_0})x(t) + \frac{(\mu_{\Theta_0}^2 - \mu_{\Theta_1}^2)}{2}t\end{aligned}\tag{13}$$

Thus both $L_\theta(x(t))$ and $Z_l(t)$ are linear in $x(t)$. In particular we assume the slopes to be symmetrical, i.e. $\mu_{\Theta_1} = -\mu_{\Theta_0} > 0$ (otherwise we can adjust the baseline to be $(\mu_{\Theta_0} + \mu_{\Theta_1})t/2$ and make the slopes symmetrical), then $Z_l(t) = 2\mu_{\Theta_1}x(t)$ is directly proportional to $x(t)$.

A.2 Bayesian inference for heterogeneous visual search

The target-present likelihood is given by marginalizing out the target location l_T , distractor mixture ϕ , as well as the target and distractor orientations:

$$\begin{aligned}
P(X|C=1) &= \sum_{l_T, \phi} P(X|l_T, \phi, C=1)P(\phi)P(l_T|C=1) \\
&= \sum_{l_T} P(l_T|C=1) \sum_{\phi} P(\phi) \sum_{\vec{Y}=\{Y_1, \dots, Y_M\}} P(X|\vec{Y})P(\vec{Y}|l_T, \phi, C=1) \\
&= \sum_{l_T} P(l_T|C=1) \sum_{\phi} P(\phi) \sum_{\vec{Y}} \prod_l (P(X_l|Y_l)P(Y_l|l_T, \phi, C=1)) \\
&= \sum_{l_T} P(l_T|C=1) \sum_{\phi} P(\phi) \prod_l \sum_{Y_l} (P(X_l|Y_l)P(Y_l|l_T, \phi, C=1)) \\
&= \sum_{l_T} P(l_T|C=1) \sum_{\phi} P(\phi) P(X_{l_T}|C_{l_T}=1) \prod_{l \neq l_T} P(X_l|\phi, C_l=0) \\
&= \sum_{l_T} P(l_T|C=1) \sum_{\phi} P(\phi) \frac{P(X_{l_T}|C_{l_T}=1)}{P(X_{l_T}|\phi, C_{l_T}=0)} \prod_i P(X_i|\phi, C_i=0)
\end{aligned}$$

where

$$\begin{aligned}
P(X_l|C_l=1) &= \sum_{\theta \in \Theta_0} P(X_l|Y_l=\theta)P(\theta|C_l=1) \\
P(X_l|\phi, C_l=0) &= \sum_{\theta \in \Theta_0} P(X_l|Y_l=\theta)\phi_{\theta}
\end{aligned}$$

Similarly, the target-absent likelihood is:

$$P(X|C=0) = \sum_{\phi} P(\phi) \prod_l P(X_l|\phi, C_l=0)$$

Define:

$$P(\phi|X) = \frac{P(\phi) \prod_l P(X_l|\phi, C_l=0)}{\sum_{\phi'} P(\phi') \prod_l P(X_l|\phi', C_l=0)}$$

Then the log likelihood ratio is:

$$\begin{aligned}
\log \frac{P(X|C=1)}{P(X|C=0)} &= \log \int_l P(l_T=l|C=1)P(X_l|C_l=1) \int_{\phi} \frac{P(\phi|X)}{P(X_l|\phi, C_l=0)} \\
&= \mathcal{S}_{l=1, \dots, M}^{\max} \left(\log P(l_T=l|C=1) + \log P(X_l|C_l=1) + \mathcal{S}_{\phi \in \Phi}^{\max} (\log P(\phi|X) - \log P(X_l|\phi, C_l=0)) \right)
\end{aligned}$$

assuming uniform prior on target locations $P(l_T=l|C=1)$ and on target type $P(Y_l=\theta|C_l=1)$:

$$\log \frac{P(X|C=1)}{P(X|C=0)} = \mathcal{S}_{l=1, \dots, M}^{\max} \left(\mathcal{S}_{\theta \in \Theta_1}^{\max} (L_{\Theta_1}(X_l)) + \mathcal{S}_{\phi \in \Phi}^{\max} \left(-\mathcal{S}_{\theta \in \Theta_0}^{\max} (L_{\theta}(X_l) + \log \phi_{\theta}) + \log P(\phi|X) \right) \right) - \log(M|\Theta_1|)$$

which proves Equation 3,4 and 5.

A.3 State formulation in visual search

We have chosen the log posterior ratios at all locations: $\vec{Z} : Z_l(t) = \log \frac{P(X_l|C_l=1)}{P(X_l|C_l=0)}, l = 1 \dots M$, to be the state of our model because the resultant system is Markov: i.e. \vec{Z} is a sufficient statistic to compute both the overall log likelihood ratio $S_{\text{homo-search}}$ and likelihood of future observations.

First, as shown in [5, 6]

$$S_{\text{homo-search}} = \log \frac{P(C=1|X)}{P(C=0|X)} = \sum_{l=1 \dots M} \mathcal{S}_{\text{max}}(Z_l) - \log(M)$$

Second, let $\Delta X = X(t+1) - X(t)$ denotes new observations at all locations at time $t+1$, the likelihood of ΔX is obtained by marginalizing the target location l_T . Denote $l_T = 0$ the target-absent event:

$$P(l_T = 0|X(t)) = P(C = 0|X) = \frac{1}{1 + \exp(S_{\text{homo-search}})} = \frac{1}{1 + \sum_l \exp(Z_l)/M}$$

$$P(l_T, l_T > 0|X(t)) = \frac{\exp(Z_{l_T})/M}{1 + \sum_l \exp(Z_l)/M}$$

For notational convenience, define $Z_0 = \log(M)$, then the equations above simplifies to:

$$P(l_T|X) = \frac{\exp(Z_{l_T})}{\sum_{l=0}^M \exp(Z_l)}$$

The posterior on l_T is sufficient to compute likelihood of ΔX :

$$P(\Delta X|X(t)) = P(\Delta X, C = 0|X(t)) + P(\Delta X, C = 1|X(t))$$

$$\text{where } P(\Delta X, C = 0|X(t)) = P(\Delta X|C = 0)P(C = 0|X(t)) = P(l_T = 0|X(t)) \prod_l P(\Delta X_l|C_l = 0)$$

$$P(\Delta X, C = 1|X(t)) = \sum_{l_T} P(\Delta X|l_T)P(l_T|X(t))$$

$$= \sum_{l_T} P(\Delta X_{l_T}|C_{l_T} = 1) \prod_{l \neq l_T} P(\Delta X_l|C_l = 0)P(l_T|X(t))$$

A.4 Translating optimal thresholds for discrimination to asymptotic thresholds for search

We discuss how to design thresholds for visual search that asymptotically achieves the best ER vs RT trade-off (as in Conjecture. 1 and Equation. 8 and 9). This is done by relating the asymptotically optimal visual search thresholds $\{\tau_-^{vs}, \tau_+^{vs}\}$ to two other pairs of thresholds:

- $\{\tau_-, \tau_+\}$: the optimal thresholds for discrimination with an *even* prior ratio (i.e. $P(C=1)/P(C=0) = 1$)
- $\{\tau'_-, \tau'_+\}$: the optimal thresholds for discrimination with a *biased* prior ratio of $1/M$.

(I) $\{\tau_-^{vs}, \tau_+^{vs}\} = \{\tau'_-, \tau'_+\}$: the asymptotic search thresholds are identical to the discrimination threshold with a $1/M$ prior ratio. The asymptotic case is where the locations $l \neq l^*$ are absolutely sure that they do not contain any target, i.e. $Z_l(t) \rightarrow -\infty, \forall l \neq l^*$. Asymptotically (i.e. after collecting a significant amount of information) this always happens when the target is absent, and happens with probability $1/M$ when the target is present (when l^* is the target location). Therefore, the asymptotic search problem can be reduced to a visual discrimination problem with a prior ratio of $1/M$.

(II) $\{\tau'_-, \tau'_+\} + \log(1/M) = \{\tau_-, \tau_+\}$: log prior ratio causes an additive change to the optimal discrimination thresholds. Let γ_+ and $-\gamma_-$ (note that $\gamma_+, \gamma_- > 0$) be the upper and lower threshold for visual discrimination with a prior of p for target-present. Let RT_C and ER_C be the expected response time and error rate when the stimulus type is $C \in \{0, 1\}$. The error rates, assuming the two thresholds are far apart, are given by (see summary in [7]):

$$\begin{aligned}\text{RT}_1(\gamma_+, \gamma_-) &\approx \text{RT}_1(\gamma_+) = kC_{\text{time}}\gamma_+ \\ \text{RT}_0(\gamma_+, \gamma_-) &\approx \text{RT}_0(\gamma_-) = kC_{\text{time}}\gamma_- \\ \text{ER}_1(\gamma_+, \gamma_-) &\approx \text{ER}_1(\gamma_-) = \frac{1}{1 + e^{\gamma_-}} \\ \text{ER}_0(\gamma_+, \gamma_-) &\approx \text{ER}_0(\gamma_+) = \frac{1}{1 + e^{\gamma_+}}\end{aligned}$$

where k is an unknown constant that is inversely proportional to the drift-rate. The total risk $\mathcal{R}(\gamma_+, \gamma_-)$ is given by:

$$\mathcal{R}(\gamma_+, \gamma_-) = p\text{RT}_1(\gamma_+) + (1-p)\text{RT}_0(\gamma_-) + p\text{ER}_1(\gamma_-) + (1-p)\text{ER}_0(\gamma_+)$$

At the optimal thresholds γ_+^* and γ_-^* , it must be that the local derivatives of the risk function w.r.t. the thresholds are zero:

$$\begin{aligned}\frac{\partial \mathcal{R}}{\partial \gamma_+} \Big|_{\gamma_+ = \gamma_+^*} = 0 &\implies kC_{\text{time}} = \frac{(1-p)e^{-\gamma_+^*}}{p(1 + e^{-\gamma_+^*})^2} \approx \frac{1-p}{p}e^{-\gamma_+^*} = e^{-(\gamma_+^* + \log \frac{p}{1-p})} \\ &\implies \gamma_+^*(p) = -\log(kC_{\text{time}}) - \log \frac{p}{1-p} \\ \frac{\partial \mathcal{R}}{\partial \gamma_-} \Big|_{\gamma_- = \gamma_-^*} = 0 &\implies \gamma_-^*(p) = -\log(kC_{\text{time}}) + \log \frac{p}{1-p}\end{aligned}$$

Setting $p = 1/2$ (or equivalently, $\log \frac{p}{1-p} = 0$) and $p = 1/(1+M)$ (or equivalently, $\log \frac{p}{1-p} = -\log(M)$) respectively, we have:

$$\begin{aligned}\tau_+ &= \gamma_+^*\left(\frac{1}{2}\right) = -\log(kC_{\text{time}}) \\ \tau'_+ &= \gamma_+^*\left(\frac{1}{1+M}\right) = -\log(kC_{\text{time}}) + \log(M) \\ \implies \tau'_+ &= \tau_+ + \log(M)\end{aligned}$$

Similarly,

$$\implies \tau'_- = -\gamma_-^*\left(\frac{1}{1+M}\right) = -(\gamma_-^*\left(\frac{1}{2}\right) - \log(M)) = \tau_- + \log(M)$$

Therefore, the optimal thresholds $\{\tau'_-, \tau'_+\}$ with a biased prior ratio may be obtained by offsetting the optimal thresholds $\{\tau_-, \tau_+\}$ with the log prior ratio.

Combining (I) and (II), see that the asymptotic visual search thresholds are given by $\{\tau_-^{vs}, \tau_+^{vs}\} = \{\tau_-, \tau_+\} + \log(M)$.

A.5 Dynamic programming

We use dynamic programming (Eq. 7) to solve for the optimal decision strategy for each (\vec{Z}, t) pair. We set up the recurrence relationship in Equation. 7 by computing for each \vec{Z} :

$$P_1(\vec{Z}) \triangleq P(C = 0|\vec{Z}) = \frac{1}{1 + \sum_{l=1}^M \exp(Z_l)/M} \quad (14)$$

$$\begin{aligned} P(\vec{Z}(t+1)|\vec{Z}(t)) &= \sum_{l_T=0}^M P(\vec{Z}(t+1)|l_T, \vec{Z}(t))P(l_T|\vec{Z}(t)) \\ &= \sum_{l_T=0}^M \mathcal{N}(\vec{Z}(t+1)|\vec{Z}(t) + \Delta\vec{Z}^{(l_T)}, V^{(l_T)}) \frac{\exp(Z_{l_T})}{\sum_{l=0}^M \exp(Z_l)} \end{aligned} \quad (15)$$

where $\Delta\vec{Z}^{(l_T)}$ and $V^{(l_T)}$ are the mean and variance of the change in \vec{Z} when the target location is l_T (Eq. 13). Let μ_0 and μ_1 be the diffusion slopes (Eq. 12) when the stimulus is a distractor and a target, respectively:

$$\begin{aligned} \Delta Z_l^{(l_T)} &= \mu_0(\mu_1 - \mu_0)\delta t & \forall l \neq l_T \\ \Delta Z_{l_T}^{(l_T)} &= \mu_1(\mu_1 - \mu_0)\delta t \\ V_l^{(l_T)} &= (\mu_{\Theta_1} - \mu_{\Theta_0})^2 \delta t & \forall l = 1, \dots, M \end{aligned}$$

Having readily computed $P_1(\vec{Z})$ and $P(\vec{Z}(t+1)|\vec{Z}(t))$, we use backward induction to solve the dynamic programming equation (Eq. 7) by starting from an infinite horizon $t = \infty$ (see below for practical details), at which point the ideal observer is forced to declare either target-present or target-absent.

In our experiments we can only handle two display locations ($\vec{Z} = \{Z_1, Z_2\}$). We discretize each Z_l from -10 to 10 into 2000 bins and time into bins of 10 ms. These parameter values are found by searching in log scale the tightest range and the most economic discretization level that do not affect the result. To simulate the infinite horizon we start backward induction from $t = T_{max}$, and increase T_{max} on a log scale until the solution for the first 10 seconds stabilizes. We find that the optimal decision strategy is constant in time, which is not surprising since the transition probabilities $P(\vec{Z}(t+1)|\vec{Z})$ and the terminal values $P_1(\vec{Z})$ are constant in time. Thus we can stop the induction whenever the decision boundaries converge in time. Typically $T_{max} = 15$ seconds is sufficient to handle diffusion slopes as low as $\pm 5/sec$.